Bayesian Optimisation of Multi-Component Promotion: A Case Study on Pt/Al_2O_3 Catalysts for Perhydro Benzyltoluene Dehydrogenation

Jakob Biesinger[a], Torben Tiezema[a], Laura M. Helleckes[b,c], Moritz Wolf*[,a,d]

a) Karlsruhe Institute of Technology (KIT), Institute of Catalysis Research and Technology, Hermann-von-Helmholtz-Platz 1, 76344 Eggenstein-Leopoldshafen, Germany.

b) Imperial College London, Department of Chemical Engineering, Exhibition Road, London SW7 2AZ, United Kingdom

c) Imperial College London, I-X Centre for AI in Science, 80-92 Wood Lane, London W12 0BZ, United Kingdom

d) Karlsruhe Institute of Technology (KIT), Engler-Bunte-Institut, Engler-Bunte-Ring 1, 76131 Karlsruhe, Germany.

*Corresponding author. E-mail address: moritz.wolf@kit.edu

Abstract

The rational design of multi-promoter catalysts is challenged by the large combinatorial design space and the high experimental costs associated with catalyst synthesis and testing. Bayesian optimisation (BO) has emerged as a promising machine-learning framework for data-efficient optimisation of expensive *black-box* problems. BO provides immense potential for accelerating discovery in catalyst promotion due to the combinatorial explosion of vast multi-variable spaces, in particular for multi-component promotion. In this work, BO is introduced as a practical decision-support tool for the optimisation of promoter compositions of a 0.30 wt.% Pt/Al_2O_3 catalyst for the dehydrogenation of the liquid organic hydrogen carrier (LOHC) perhydro benzyltoluene. Beyond demonstrating the practical application of BO, the presented framework is systematically analysed with respect to surrogate model selection and acquisition strategy, providing general guidance for its application in catalyst optimisation. Conventional analysis of an initial experimental dataset reveals characteristic trends for individual promoter effects on catalyst activity and selectivity, which highlights the complex trade-offs governing catalyst design. Experimental optimisation campaigns demonstrate that BO efficiently explores the constrained catalyst design space and identifies chemically meaningful catalyst compositions within only a few optimisation iterations. Overall, this study demonstrates the applicability of BO as a data-efficient framework for catalyst development and provides a foundation for future optimisation

campaigns incorporating larger datasets, additional design variables, and advanced constraint-aware multi-objective strategies.

1. Introduction

1.1. Catalyst Promotion for Perhydro Benzyltoluene Dehydrogenation

Liquid organic hydrogen carriers (LOHCs) have emerged as a promising technology for chemical storage of hydrogen. Among the various LOHC systems reported in literature, perhydro benzyltoluene (H12-BT) has attracted considerable attention due to its high hydrogen storage capacity, favourable thermodynamic properties, chemical stability, and compatibility with existing liquid fuel infrastructure facilitating global transportation.[1] Hydrogen release from H12-BT yielding benzyltoluene (H0-BT) is achieved *via* catalytic dehydrogenation (Figure 1), which represents a key process step governing the overall efficiency of the LOHC technology. Besides catalytic activity, selectivity is crucial in H12-BT dehydrogenation, as undesired consecutive reactions can lead to polycyclic aromatic by-products, such as methylfluorene species (MeFl). The high melting point of these species can impair handling of the LOHC mixture and, more importantly, their formation may promote catalyst deactivation. Hence, suppression of the by-product MeFl is essential for efficient hydrogen release.[2]

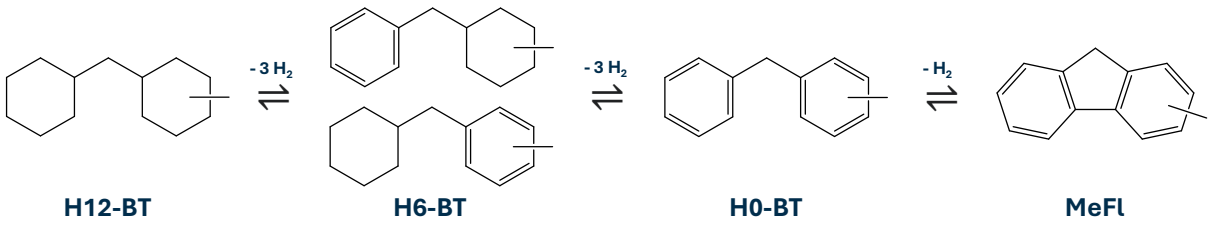


Figure 1: Simplified reaction pathway for the dehydrogenation of H12-BT via H6-BT to H0-BT and the associated consecutive formation of the undesired by-product methylfluorene (MeFl).[1]

Pt/Al_2O_3 based catalysts represent the state-of-the-art catalyst for H12-BT dehydrogenation, which typically proceeds *via* the stable intermediates of partially hydrogenated benzyltoluene (H6-BT) (Figure 1).[3] Under industrially relevant reaction conditions, Pt/Al_2O_3 catalysts only exhibit moderate catalytic performance, as evidenced by low intrinsic activity and formation of by-products.[1] Catalyst promotion has therefore emerged as a widely applied strategy to enhance catalyst performance in heterogeneous catalysis.[4–6] Promoters can modify textural, structural, and electronic properties of the active phase or the support material, thereby affecting catalyst activity, activation behaviour, selectivity, or stability.[7] For hydrocarbon dehydrogenation, a range of promoters have been reported to influence activity and stability. Table 1 summarises representative promoted catalysts, with particular emphasis on LOHC candidates, such as H12-BT, perhydro dibenzyltoluene (H18-DBT), or methylcyclohexane (MCH). The reported studies demonstrate that promoter addition can enhance activity, selectivity, or both. However, the observed effects strongly depend on the system itself, promoter identity, loading, and catalyst composition, while no superior promoter has been identified. Sulphur promotion has been commercially established *via* selective poisoning of Pt for suppression of the formation of by-products during LOHC dehydrogenation.[8]

Table 1: Literature examples of promoter species reported for Pt-based catalysts used in the dehydrogenation of hydrocarbons, using supported Pt-based catalysts under realistic experimental conditions, with a focus on LOHC systems. The table summarises observed effects on catalytic activity and selectivity, where selectivity refers to the suppression of undesired side reactions and by-product formation.

Promoter	System	Support	Activity	Selectivity	Reference
Cr	MCH	γ-Al_2O_3	▲	▲	9
Ge	n-Butane	γ-Al_2O_3	▲	▼	10
Mn	H12-BT	γ-Al_2O_3	▲	▼	11
Mo	H12-BT	γ-Al_2O_3	▲	▲	11
Mo	MCH	γ-Al_2O_3	▲	▲	9
Na	MCH	γ-Al_2O_3	▼	▲	9
Ni	H12-BT	γ-Al_2O_3	-	▼	12
Re	H12-BT	γ-Al_2O_3	▲	▼	13
S	H12-BT	a-TiO_2	▼	▲	14
S	H18-DBT	Al_2O_3	▲	▲	15
Sn	n-Decane	$ZnAl_2O_4$	▲	-	16
Sn	n-Butane	γ-Al_2O_3	▲	▲	10
Sn	Propane	θ-Al_2O_3, ZSM-5 γ-Al_2O_3	▲	▲	6, 17, 18
V	Propane	SiO_2	-	▲	19
V	MCH	γ-Al_2O_3	▲	▲	9
W	H18-DBT	γ-Al_2O_3	▲	▼	20

While many studies systematically focus on the effect of individual promoters, a single promoter is often not able to improve all relevant performance criteria simultaneously. In practice, several promoter species are frequently combined to balance competing effects, such as enhancing activity while maintaining selectivity or suppressing catalyst deactivation.[21] Hence, catalyst promotion results in a high-dimensional catalyst design space, in which several compositional parameters must be adjusted. Exploring such a design space experimentally is challenging, as each catalyst synthesis and performance evaluation is time-consuming and resource-intensive. Moreover, the relationship between catalyst composition and catalytic performance is generally complex, highly nonlinear, and not accessible through analytical models.

As a result, catalyst optimisation for H12-BT dehydrogenation constitutes a prototypical example of an expensive-to-evaluate optimisation problem, where experimental insight can only be obtained through sequential extensive laboratory experiments. This motivates the application of data-efficient optimisation strategies, capable of guiding

experimental campaigns towards improved catalyst candidates with a minimal number of experiments. BO is ideally suited for such a research question and has been recently applied for the optimisation of multi-component catalyst systems.[6,22–27]

1.2. Bayesian Optimisation for Expensive-to-Evaluate Problem Statements

BO is a machine learning based global optimisation framework for continuous *black-box* objective functions $f(\mathbf{x})$. In this context, *black-box* refers to functions for which neither an analytical expression nor derivative information is available, and where function evaluations may be (computationally) expensive or time-consuming. The problem statement for the optimisation can be described as

$$\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}), \tag{1}$$

where $\mathcal{X} \subseteq \mathbb{R}^n$ denotes the feasible search space of the n-dimensional parameter vector $\mathbf{x}$. In the context of heterogeneous catalysis, $\mathbf{x}$ may represent catalyst design variables such as active metal or promoter identity or loading, or choice of support material, while $f(\mathbf{x})$ represent catalyst performance metrics, such as educt conversion or selectivity towards a desired product, which depend on these design variables. BO is particularly well suited for such expensive-to-evaluate functions, as it aims to identify the global optimum of f using a minimal number of function evaluations. To this end, BO adaptively selects new evaluation points by balancing *exploitation* of regions with high observed objective values and *exploration* of regions associated with high predictive uncertainty.

In its standard formulation, a BO framework consists of two main components. Firstly, a probabilistic surrogate model that approximates the unknown objective function, and secondly, an acquisition function that guides the selection of new query points based on the surrogate's predictions and associated uncertainties (Figure 2a). BO proceeds in a sequential manner, where experimental results are incorporated into the surrogate model to iteratively refine future experimental decisions, thereby creating a feedback loop between modelling and experimentation. Both components are described in more detail in the following.

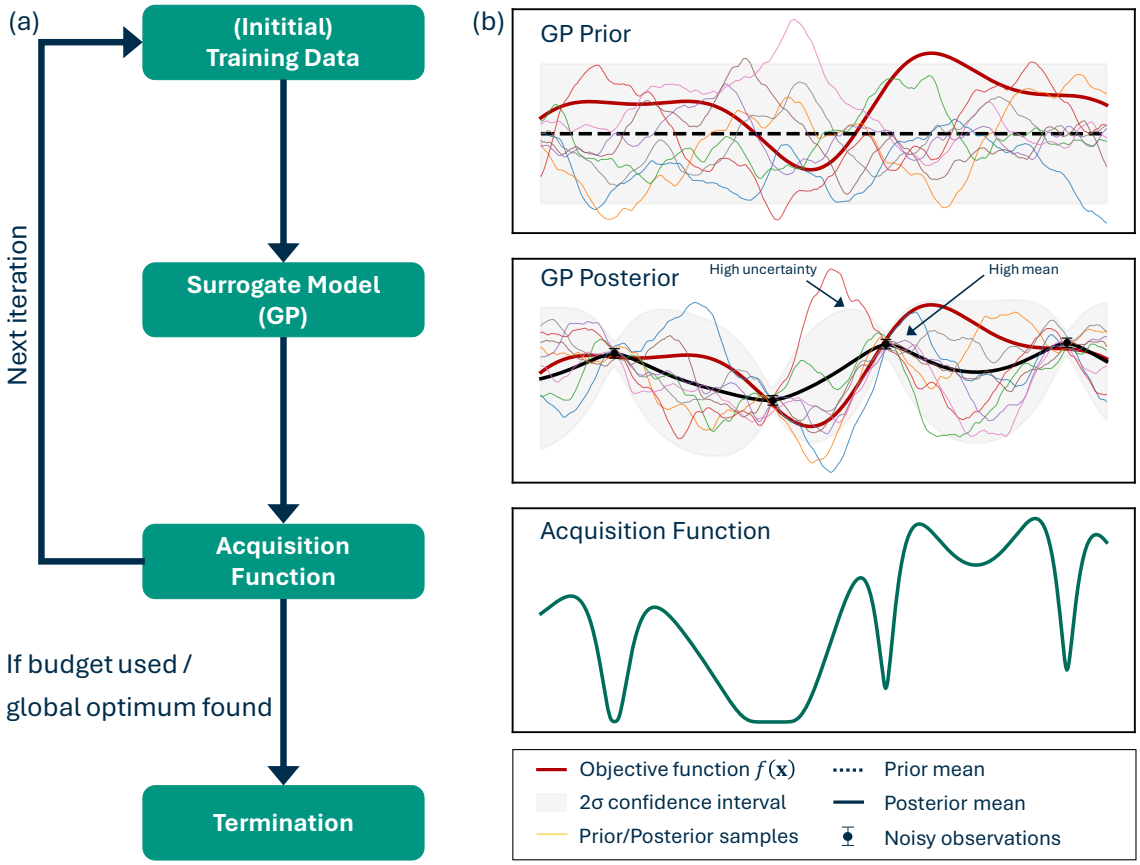


Figure 2: (a) Schematic illustration of the BO workflow. (b) A Gaussian process (GP) prior is conditioned to the available data and used to construct an acquisition function (here: Expected Improvement) that balances exploration and exploitation. The acquisition function proposes new candidate point(s), which are experimentally evaluated, and the resulting data are iteratively incorporated into the model in a closed-loop manner.

Surrogate models provide a probabilistic approximation of the unknown objective function $f(\mathbf{x})$. Rather than exactly reproducing the underlying function, the surrogate is primarily required to capture sufficient structural information to identify promising regions of the design space. Examples of surrogate models are Gaussian processes (GPs), neural networks or random forests, with GPs typically being used in BO due to their analytical tractability, inherent uncertainty quantification, and their suitability for small datasets, as commonly encountered in expensive experimental applications.

GPs are Bayesian models that define distributions over functions (Figure 2b). In a Bayesian setting, a GP is used as a prior distribution over the objective function, which is updated to a posterior distribution once experimental data are observed. Formally, a GP can be written as

$$f(\mathbf{x}) \sim \mathcal{GP}\big(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')\big), \tag{2}$$

where $m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})]$ is the mean function (prior mean) and $k(\mathbf{x}, \mathbf{x}') = \mathrm{Cov}[f(\mathbf{x}), f(\mathbf{x}')]$ is the kernel (covariance) function, which specifies the covariance between function values at two input locations in the design space $\mathbf{x}$ and $\mathbf{x}'$. The kernel encodes prior assumptions about the structural properties of the objective function. Various kernels differ in terms of

the *a priori* structural properties of a function they express, such as smoothness or variance.

A commonly used kernel that produces infinitely differentiable (very smooth) functions is the Radial Basis Function (RBF) (or Squared Exponential) kernel, given by

$$k_{\mathrm{RBF}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \cdot \exp\left(\frac{-\| \mathbf{x} - \mathbf{x}' \|^2}{2\ell^2}\right), \tag{3}$$

where σ_f^2 denotes the signal variance, ℓ the characteristic length scale and $\|\bullet\|$ the Euclidean distance.[28] The kernel assumes that design variables close to each other, such as those representing similar catalyst compositions, are associated with similar objective values, with the correlation decreasing as their distance increases. The characteristic length scale determines how rapidly this correlation decays, where small values permit rapid local variations, while larger values enforce smoother changes of the objective function.

The Rational Quadratic (RQ) kernel can be interpreted as a scale mixture of RBF kernels with different length scales and is defined as

$$k_{\mathrm{RQ}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \cdot \left(1 + \frac{\| \mathbf{x} - \mathbf{x}' \|^2}{2\alpha\ell^2}\right)^{-\alpha}, \tag{4}$$

where the parameter $\alpha > 0$ controls the relative weighting of large- and small-scale variations.[28] In the limit $\alpha \to \infty$, the RQ kernel converges to the RBF kernel.

Matérn kernels are a generalisation of the RBF kernel and only allow for limited differentiability. Therefore, these kernels are often better suited for less smooth or noisy experimental data. The Matérn kernel is given by

$$k_{\mathrm{Matérn\ }\nu}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \cdot \frac{2^{1-\nu}}{\Gamma(\nu)} \cdot \left(\frac{\sqrt{2\nu} \cdot \| \mathbf{x} - \mathbf{x}' \|}{\ell}\right)^{\nu} \cdot K_\nu\left(\frac{\sqrt{2\nu} \cdot \| \mathbf{x} - \mathbf{x}' \|}{\ell}\right), \tag{5}$$

where ℓ, ν and σ are kernel hyperparameters, $\Gamma(\cdot)$ denotes the Gamma function, and $K_\nu(\cdot)$ is the modified Bessel function of the second kind.[28] The smoothness parameter ν directly controls the differentiability of functions drawn from the GP, making a Matérn GP $\lfloor \nu \rfloor$ times mean-square differentiable. Common choices are ν = 3/2 and ν = 5/2, which correspond to once and twice differentiable functions, respectively.[28] In GP regression, kernel hyperparameters are typically determined by maximising the marginal log-likelihood of the observed data, thereby adapting the covariance structure of the model to the dataset.[28]

Given the learned hyperparameters, GP regression is performed by conditioning the prior on the observed training data $\mathcal{D}$ with noisy experimental observations y_i

$$\mathcal{D} = \{(\mathbf{x}_i, \mathrm{y}_i)\}_{i=1}^N, \qquad \mathrm{y}_i = f(\mathbf{x}_i) + \varepsilon_i, \qquad \varepsilon_i \sim \mathcal{N}(0, \sigma_n^2), \tag{6}$$

where ε_i denotes independent observation noise normally distributed with zero mean and noise variance σ_n^2. Applying Bayes' theorem yields the predictive posterior distribution

$$p(f(\mathbf{x}) \mid \mathcal{D}) \sim \mathcal{N}\big(\mu(\mathbf{x}), \sigma^2(\mathbf{x})\big), \tag{7}$$

which provides both a mean prediction $\mu(\mathbf{x})$ and a corresponding uncertainty estimate $\sigma(\mathbf{x})$ at any input location $\mathbf{x}$ of the design space. As the posterior predictive distribution is Gaussian, $\sigma(\mathbf{x})$ quantifies the posterior uncertainty, such that approximately 68% and 95% credible intervals correspond to $\mu(\mathbf{x}) \pm \sigma(\mathbf{x})$ and $\mu(\mathbf{x}) \pm 2\sigma(\mathbf{x})$, respectively. This quantified uncertainty is a key property enabling the effective exploration-exploitation trade-off in BO.

The second fundamental component of BO is the acquisition function $a(\mathbf{x})$, which determines the locations of the design space where the objective function is evaluated next (Figure 2). This function is the means to realising a certain optimisation policy, i.e., the policy with which new experiments are selected. While the simplest policies could, for example, randomly select new locations, most acquisition functions balance exploration and exploitation in practice. They are constructed using the predictive posterior distribution of the surrogate model and therefore explicitly depend on both the predicted mean and the associated uncertainty. At each iteration, the next query point $\mathbf{x}_{\text{next}}$ is obtained by maximising the acquisition function

$$\mathbf{x}_{\text{next}} = \arg\max_{\mathbf{x} \in \mathcal{X}} a(\mathbf{x} \mid \mathcal{D}_N), \tag{8}$$

where $\mathcal{D}_N$ denotes the set of N observations collected so far.

A widely used acquisition function is the Expected Improvement (EI)[29], which aims to maximise the expected improvement over the current best observed function value f^*. It is defined as

$$\text{EI}(\mathbf{x}) = \mathbb{E}[\max(0, f(\mathbf{x}) - f^*) \mid \mathcal{D}_N]. \tag{9}$$

For the special case of noiseless single-objective BO, EI admits a closed-form expression under a GP posterior. This is given by

$$\text{EI}(\mathbf{x}) = (\mu(\mathbf{x}) - f^*)\Phi(Z) + \sigma(\mathbf{x})\phi(Z), \qquad Z = \frac{\mu(\mathbf{x}) - f^*}{\sigma(\mathbf{x})}, \tag{10}$$

where $\Phi(\cdot)$ and $\phi(\cdot)$ denote the cumulative distribution function and probability density function of the standard normal distribution, respectively. However, in more general settings involving batch evaluations, observation noise, or multiple objectives, the acquisition function is defined as a high-dimensional expectation that is not analytically tractable. In such cases, EI is efficiently approximated, for example using quasi-Monte-Carlo integration based on Sobol sequences.[30,31]

Another popular acquisition function is the Upper Confidence Bound (UCB)[32], which explicitly trades off exploitation and exploration by augmenting the posterior mean with a multiple of the predictive standard deviation. UCB is given by

$$\mathrm{UCB}(\mathbf{x}) = \mu(\mathbf{x}) + \lambda \cdot \sigma(\mathbf{x}), \tag{11}$$

where the parameter $\lambda > 0$ controls the exploration-exploitation trade-off.

1.3. Practical Extensions of Bayesian Optimisation

Beyond the aforementioned standard formulation, several extensions to BO have been developed to address practical experimental settings. In many real-world problems, including catalyst optimisation, multiple objectives must be considered simultaneously, such as catalyst activity, selectivity, and stability. These objectives are often inherently conflicting, where highly active catalysts may also promote undesired side reactions leading to reduced selectivity. In such cases, multi-objective (MO) BO aims to identify the Pareto front, i.e., the set of non-dominated candidates representing optimal trade-offs between the objectives. A candidate is Pareto-optimal if no other candidate exists that is at least as good in all objectives and strictly better in at least one objective. While this provides a set of optimal trade-off candidates and allows for exploring between competing objectives, selecting a final design from the Pareto front typically requires additional decision-making or weighting of the objectives. In contrast, single-objective (SO) optimisation considers a single objective function and therefore directly yields a single optimal solution.

Furthermore, BO can be adapted to propose multiple candidates in parallel through batch optimisation, which is particularly relevant for experimental workflows where several experiments can be conducted simultaneously. However, selecting multiple candidates is not trivial, as evaluating one candidate changes the information available for subsequent selections. Batch BO therefore requires specialised strategies to approximate this sequential decision process when proposing several candidates at once. Finally, constraint-aware BO incorporates feasibility requirements into the optimisation process by restricting the search to feasible regions of the design space, for example through limits on catalyst composition or synthesis constraints.

To this point, we summarised the central concepts of BO. Readers interested in a more detailed treatment are referred to standard references in the field, such as the comprehensive introduction by R. Garnett[33].

1.4. Aim of the Study

Herein, the expensive-to-evaluate multi-component promotion of a 0.30 wt.% Pt/Al_2O_3 catalyst for dehydrogenation of the LOHC H12-BT will be tackled by single- and multi-objective BO using a training dataset with literature-inspired catalyst compositions to maximise catalyst activity (SO) and selectivity, while minimising catalyst costs (MO). The training data set provides initial information on the performance of 38 catalyst candidates with various promoter combinations, while trends are extracted in a conventional correlative analysis. For BO, various GP kernels and acquisition function are compared to provide a suitable framework for catalyst optimisation. Rather than aiming at an

exhaustive optimisation of the catalyst system, this study primarily serves to demonstrate the practical applicability of BO to the optimisation of multi-component promoted catalysts using simple BO optimisation frameworks and to provide a methodological basis for future, more extensive optimisation studies.

2. Materials and Methods

2.1. Catalyst Preparation and Characterisation

Boehmite (PURAL TH130, Sasol Germany) was calcined in a muffle furnace at 1000 °C for 3 h yielding a θ-δ-Al_2O_3 support material, as evidenced by X-ray diffraction analysis. The support material was further characterised by N_2-physisorption. Further details can be found in section S1.1 of the Supporting Information. Promoted Pt/θ-δ-Al_2O_3 catalysts were prepared *via* wet co-impregnation using the corresponding metal precursors (Table2, Table S1).[13] The required amounts of Pt and promoter precursor were dissolved in Milli-Q water and the θ-δ-Al_2O_3 material was added under stirring targeting a variety of promotor loadings and combinations (Table S2). The catalysts were labelled as ω_{Pt}Pt-$\omega_{x,i}X_i$, with ω_{Pt} denoting Pt target loading and $\omega_{x,i}$ the target loading of promotor(s) X_i in wt.%. Subsequently, the solvent was removed in a rotary evaporator by stepwise reduction of the pressure from 140 to 20 mbar. The dried powder was calcined in a muffle furnace at 400 °C for 4 h under ambient air using a heating rate of 10 K min^{-1}. Afterwards, the catalysts were placed in ceramic boats inside a tubular furnace and reduced at 400 °C for 2 h in 10% H_2/N_2 with a heating rate of 10 K min^{-1} yielding the activated catalysts. For selected samples (Table S2), selective poisoning with S was carried out in an additional impregnation step using ammonium sulphate as S precursor (Table 2), which was followed by the aforementioned steps. The actual metal and sulphur loadings of prepared catalysts were analysed by inductively coupled plasma atomic emission spectrometry (ICP-AES) as described in S1.1.

Table 2: Overview and details of precursor materials used for catalyst preparation.

Element	Precursor	Purity (%)	Supplier
Ba	$Ba(NO_3)_2$	≥ 99	Thermo Scientific
Ce	$Ce(NO_3)_3 \cdot 6\ H_2O$	≥ 99.5	Thermo Scientific
Cr	$Cr(NO_3)_3 \cdot 9\ H_2O$	≥ 99.99	Thermo Scientific
Mo	$(NH_4)_6Mo_7O_{24} \cdot 4\ H_2O$	≥ 99	Thermo Scientific
Na	$NaNO_3$	≥ 99.5	Merck
Pt	$H_2[PtCl_6] \cdot 6\ H_2O$	≥ 99.9	Merck
Re	Re_2O_7	≥ 99.9	Merck
S	$(NH_4)_2SO_4$	≥ 99.5	Thermo Scientific
Sn	$SnCl_4 \cdot 5\ H_2O$	≥ 98	Thermo Scientific
Sr	$Sr(NO_3)_2$	≥ 99	VWR
V	NH_4VO_3	≥ 99.5	Thermo Scientific
W	$(NH_4)_6H_2W_{12}O_{40} \cdot x\ H_2O$	≥ 90 WO_3	Carl Roth

2.2. Catalytic Dehydrogenation

Dehydrogenation of H12-BT (Hydrogenious LOHC Technologies, Germany, 99.7 % purity) was carried out in a five-neck round bottom flask equipped with an argon inlet, a thermocouple immersed in the LOHC, a catalyst spoon, a reflux condenser, and a septum enabling LOHC sampling.[13] A total of 32.2 ml H12-BT was heated to 250 °C and 1.079 g of catalyst (equals a molar LOHC:Pt ratio of 10^4) was added to the reaction mixture under continuous magnetic stirring at 600 min^{-1} and ambient pressure. A 300 ml min^{-1} flow of Ar was used as carrier gas for released hydrogen to enable inline quantification by a thermal conductivity detector (TCD, FTC 300, Messkonzept GmbH). Liquid samples were taken periodically through a neck covered with a septum using a syringe and analysed by a gas chromatograph (GC 8890, Agilent Technologies, Inc.) equipped with flame ionisation detector (FID) and a DB-17ms column (length: 30 m; inner diameter: 250 µm) yielding the composition of the liquid phase. Reaction progress was quantified by calculating the degree of dehydrogenation (DoDH) using Equation (12)

$$\mathrm{DoDH} = 1 - (x_{\mathrm{H12-BT}} + 0.5 \cdot x_{\mathrm{H6-BT}}), \tag{12}$$

where x_{i} denotes the molar fraction of the component i in the liquid phase. Time-resolved catalyst productivity was calculated using the mass flow of released hydrogen ($\dot{m}_{\mathrm{H_2}}$) from the TCD according to Equation (13)

$$\mathrm{Productivity} = \frac{\dot{m}_{\mathrm{H_2}}}{m_{\mathrm{Pt}}}, \tag{13}$$

where m_{Pt} denotes the mass of the active component Pt using the nominal Pt loading of 0.30 wt.%.

2.3. Catalyst Performance Evaluation and Constraints

Catalyst performance was evaluated using three optimisation objectives: catalyst activity, selectivity, and costs. Catalyst activity was defined as an aggregate performance metric derived from multiple experimentally measured metrics, including the time required to reach a DoDH of 0.5, the integral of the DoDH over the duration of the experiment, and the maximum productivity (Equation 13). These metrics capture complementary aspects of catalytic activity and were combined to obtain a single activity score as described below.

Selectivity was quantified based on the formation of undesired MeFl species as identified by GC-FID analysis. Specifically, the molar fraction of these species at a DoDH of 0.75 and the cumulative MeFl formation over the reaction course were used as experimentally measured metrics. Analogous, these metrics were combined to yield a single selectivity metric.

The cost metric quantifies the precursor-related material costs of a catalyst system based on the promoter nominal loadings and the corresponding precursor prices per gram of

promoter element. It is calculated as a weighted sum over all promoters, as detailed in section S1.2.

The abovementioned metrics were individually normalised to [0,1] based on the available training dataset $\mathcal{D}_{\text{start}}$ prior to catalyst optimisation and subsequently aggregated into the optimisation objectives. The optimisation aimed to maximise activity while minimising selectivity towards MeFl and catalyst costs. To obtain a consistent maximisation formulation, experimental metrics to be minimised were inverted during normalisation such that higher values correspond to more desirable objective values. The full mathematical definition is provided in the Supporting Information (section S1.2). The catalyst composition was limited to the following constraints:

- Pt loading ω_{Pt} = 0.30 wt.%.
- Max. 3 promotors out of {Ba, Ce, Cr, Mo, Na, Re, Sn, Sr, V, W} + optional {S}.
- Promoter metal loading 0.06 ≤ ω_{Me} ≤ 0.90 wt.% and Sulphur loading 0.03 ≤ ω_{S} ≤ 0.30 wt.%, if included.

2.4. Bayesian Optimisation Framework

The BO workflow was implemented in Python 3.13 using the BoTorch library (version 0.16)[34], a modern framework for BO built on PyTorch[35]. GP regression was performed using GPyTorch[36]. The surrogate model was implemented as a `ModelListGP` architecture provided by the BoTorch library[34], comprising three independent single-task GPs, one for each objective according to

$$f(\mathbf{x}) = \left(f_{\text{Activity}}(\mathbf{x}), f_{\text{Selectivity}}(\mathbf{x}), f_{\text{Costs}}(\mathbf{x})\right)^{\top}. \tag{15}$$

Kernel hyperparameters were optimised *via* marginal log-likelihood maximisation using L-BFGS-B, a standard gradient-based optimiser for GP hyperparameter estimation. Conditioning on the observed data, the GP model provides a predictive posterior over a candidate set $X \subset \mathcal{X}_{Pool}$ (see below)

$$f_i(X) \mid \mathcal{D} \sim \mathcal{N}\left(\mu_i(X), \Sigma_i(X)\right), \qquad X = \{\mathbf{x}_1, \ldots \mathbf{x}_q\}, \tag{16}$$

where the joint covariance matrix $\Sigma(X)$ is block-diagonal as the objectives are modelled independently. Catalyst design was represented as a 11-dimensional feature vector $\mathbf{x}$ including the loading of the eleven promotors with respect to the aforementioned constraints according to

$$\mathbf{x} = (\omega_{\text{Ba}}, \omega_{\text{Ce}}, \omega_{\text{Cr}}, \omega_{\text{Mo}}, \omega_{\text{Na}}, \omega_{\text{Re}}, \omega_{\text{Sn}}, \omega_{\text{Sr}}, \omega_{\text{V}}, \omega_{\text{w}}, \omega_{\text{S}})^{\top}. \tag{17}$$

In this study, we explore both SO and MO optimisation settings, as introduced above (section 1.3). For SO optimisation, we considered EI[29], Noisy-EI (NEI)[37] accounting for observation noise, and UCB[32] acquisition functions to propose catalyst composition. For MO optimisation, Expected Hypervolume Improvement (EHVI)[38] and its noise-aware

variant NEHVI acquisition functions[39] were considered, both maximising the expected hypervolume improvement, the volume of objective space dominated by the Pareto front. Furthermore, UCB with ParEGO-style (Pareto Efficient Global Optimisation) Tchebycheff scalarisation[40], which uses random weight vectors to explore the Pareto front, was implemented. As mentioned earlier, the acquisition function cannot be calculated analytically for these settings. Hence, posterior samples

$$f^{(s)}(X) \sim p(f(X)|\,\mathcal{D}) \tag{18}$$

were drawn using S Monte-Carlo (MC) samples[30] (Sobol quasi-random sequences[31]) to approximate the acquisition function expectation as

$$\hat{a}(X) = \frac{1}{S}\sum_{s=1}^{S} g(\,f^{(s)}(X), X). \tag{19}$$

A discrete candidate pool $\mathcal{X}_{Pool}$ was generated using Sobol quasi-random sequences[31] within the constrained compositional design space (2.3), which provides better space-filling properties than pseudo-random sampling. To select a batch of candidate size q, the acquisition function was evaluated pointwise for all candidates in the pool and used to rank them. Candidates were then selected in a greedy manner, starting from the highest-ranked one, while enforcing a minimum L1-distance in compositional space to ensure diversity within the batch.

To asses GP model and predictive quality and identify a suitable kernel for the subsequent optimisation, two validation strategies were applied using the available training dataset D_{Start} prior to the catalyst optimisation comparing various kernels: (i) leave-one-out cross-validation (LOOCV) to probe sensitivity to individual data points and (ii) repeated random train/test splits (50/50 and 85/15) to mimic prospective deployment with different data availability. Therefore, corresponding catalyst candidates from D_{Start} were sequentially held out, while the model was retrained on the remaining data and the held-out point(s) was/were predicted.

The performance of the aforementioned acquisition functions was analysed through retrospective (offline) optimisation. In this setting, the available dataset D_{Start} was partitioned into an initial training set and a pool of candidate points. The initial training set was used to fit the surrogate model, while the remaining data served as a discrete candidate pool from which new points were sequentially sampled. Candidate suggestions were generated by the BO framework with various strategies, and their corresponding objective values were obtained from the existing experimental measurements. Random search was included as a baseline by randomly sampling candidates from the remaining candidate pool. These experiments were repeated over multiple sequential selection steps and random seeds, each initialised with a different randomly selected subset of D_{Start}, to identify a robust acquisition strategy for the subsequent catalyst optimisation.

To further assess the implementation, a retrospective validation was performed using the catalyst optimisation study of Ramirez *et al*. as a literature case study.[23] The corresponding methodology and results are provided in S1.3.

3. Results and Discussion

3.1. Correlative Analysis of the Training Data

The initial catalyst pool (training dataset) $\mathcal{D}_{\text{start}}$ for optimisation of catalyst composition for the herein targeted H12-BT dehydrogenation consists of 38 Pt/θ-δ-Al_2O_3 catalysts using various promoters adhering the defined constraints (see section 2.3). The promoter set (Figure 3) was selected to represent a broad range of chemical classes across the periodic table, including rare earth elements (Ce), transition metals (Mo, V, Cr, W, Re), alkali metals (Na), alkaline earth metals (Ba, Sr), main-group metals (Sn), and non-metallic elements (S). The selection was guided by two complementary considerations. Firstly, promoters reported in literature to enhance activity or selectivity in the dehydrogenation of hydrocarbons (Table 1) were included to incorporate established catalytic insights. Secondly, additional promotors were selected based on their availability to ensure a systematic and unbiased exploration of the periodic table. This combined strategy enables a comparative assessment of promoter effects across chemically distinct element classes and facilitates the identification of structure-performance relationships beyond trends established in literature.

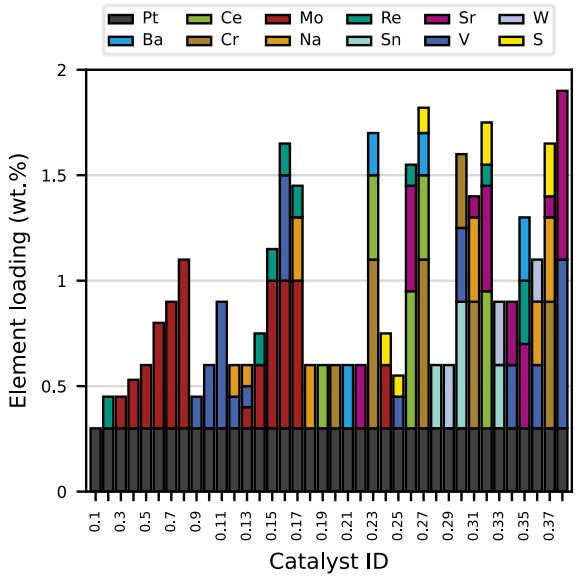


Figure 3: Compositional overview of the initial catalyst pool $\mathcal{D}_{start}$ used for model training. Each bar represents one catalyst labelled by its ID, showing the targeted elemental loadings of the active phase and individual promoter species.

While the initial dataset was not generated using a formal space-filling design, such as random, Latin hypercube, or Sobol sampling, it reflects a realistic and practical starting point for catalyst optimisation, where prior knowledge and existing laboratory data constitute the starting point for subsequent optimisation efforts. Importantly, the initial data cover all promotors and performance regimes, thereby providing a sufficiently

diverse basis for training the surrogate model in the early stages of optimisation. This approach allows for the BO framework to leverage existing experimental knowledge while iteratively improving catalyst performance through targeted exploration of the remaining parameter space. BO is particularly well suited for such settings, as it is capable of incorporating heterogeneous and non-uniformly distributed initial data.

The experimental results and evaluation of the metrics of all 38 catalysts from $\mathcal{D}_{\text{start}}$ are summarised in the SI (Table S2, Figures S8-S63). Pronounced trade-offs between activity, selectivity, and costs, especially between the former two, are identified in the 2D projections of the objective space, resulting in a well-defined Pareto front (Figure 4). Note that some of these appear dominated in the 2D projections, highlighting the importance of simultaneous consideration of all three objectives. Furthermore, the absence of a single catalyst simultaneously optimising all objectives confirms the intrinsic multi-objective nature of the optimisation problem and motivates use of a Pareto-based BO strategy. Several catalysts lie close to the Pareto front, while others are clearly dominated, indicating that the dataset covers both favourable and unfavourable performance regimes. Out of the 38 catalysts in $\mathcal{D}_{\text{start}}$, 12 were identified as Pareto-optimal, corresponding to 32% of the dataset. This relatively large fraction highlights the absence of a single dominant system. Notably, the unpromoted 0.30Pt (#0.1) reference catalyst is located on the Pareto front, indicating that promoter addition does not trivially dominate the optimisation landscape.

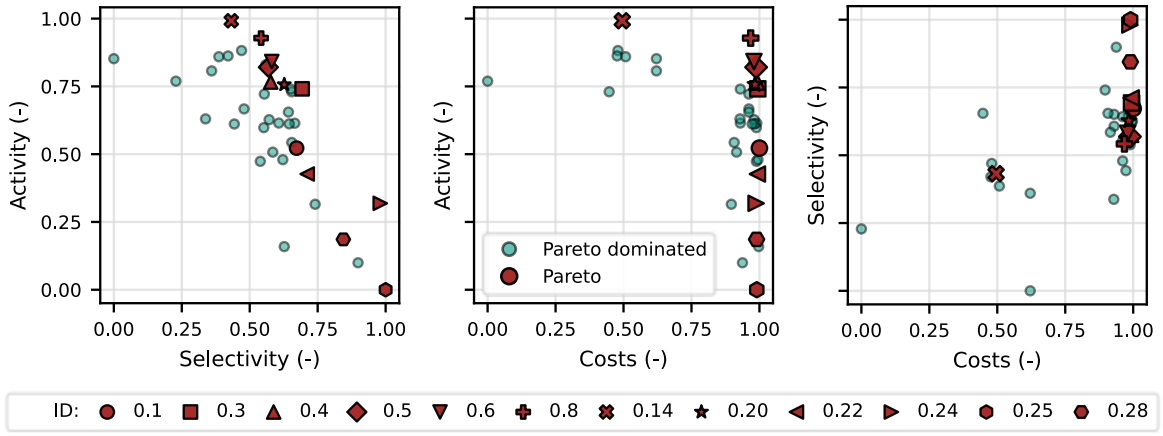


Figure 4: Two-dimensional projections of the three-dimensional objective space spanned by activity, selectivity, and cost for the initial catalyst dataset. Pareto-optimal catalysts are highlighted, while dominated catalysts are shown for reference. Note that Pareto-optimality is defined in the full three-dimensional space and may not be preserved in individual two-dimensional projections.

Among the Pareto-optimal catalysts, Mo emerges as the most prevalent promoter in 58% of Pareto-optimal candidates. At the same time, around 58% of all Mo-containing catalysts are Pareto-optimal, suggesting a favourable but non-exclusive influence of this promoter on overall catalyst performance. While Mo appears frequently among Pareto-optimal catalysts, this reflects its favourable performance across multiple objectives rather than a sampling bias. As Mo-containing catalysts are also present among

dominated solutions, use of this promoter alone is not sufficient to guarantee Pareto-optimality. Interestingly, only three Pareto-optimal catalysts contain more than one promoter. In contrast, the majority of Pareto-optimal catalysts rely on a single promoter, underscoring that increased compositional complexity does not automatically translate into superior MO performance.

The training data is further analysed by an exploratory comparison within the multivariate catalyst design space (Figure 5a) and a controlled comparison of isolated promoter effects relative to the unpromoted 0.30Pt catalyst (Figure 5b). Re clearly exhibits a pronounced increase in catalytic activity, which is accompanied by a decrease in selectivity. This trade-off is consistent with literature reports for H12-BT dehydrogenation.[13] Owing to its high precursor cost, catalysts containing Re are associated with the lowest cost objective in the exploratory analysis. Importantly, the trends observed in the multivariate comparison and the isolated Pt-Re comparison are fully consistent. Similar agreement with literature is observed for Mo, which moderately enhances activity while reducing selectivity.[11] A reverse effect is identified for Sn, which is not consistent with the frequently described activity enhancing effect in literature.[10,16–18] Notably, the trimetallic 0.30Pt-0.30Mo0.15Re catalyst (#0.14) exhibits a higher activity than all of the corresponding bimetallic systems, indicating a non-trivial synergistic effect between Re an Mo. At the same time, this catalyst exhibits the highest activity of all catalysts from $\mathcal{D}_{\text{start}}$.

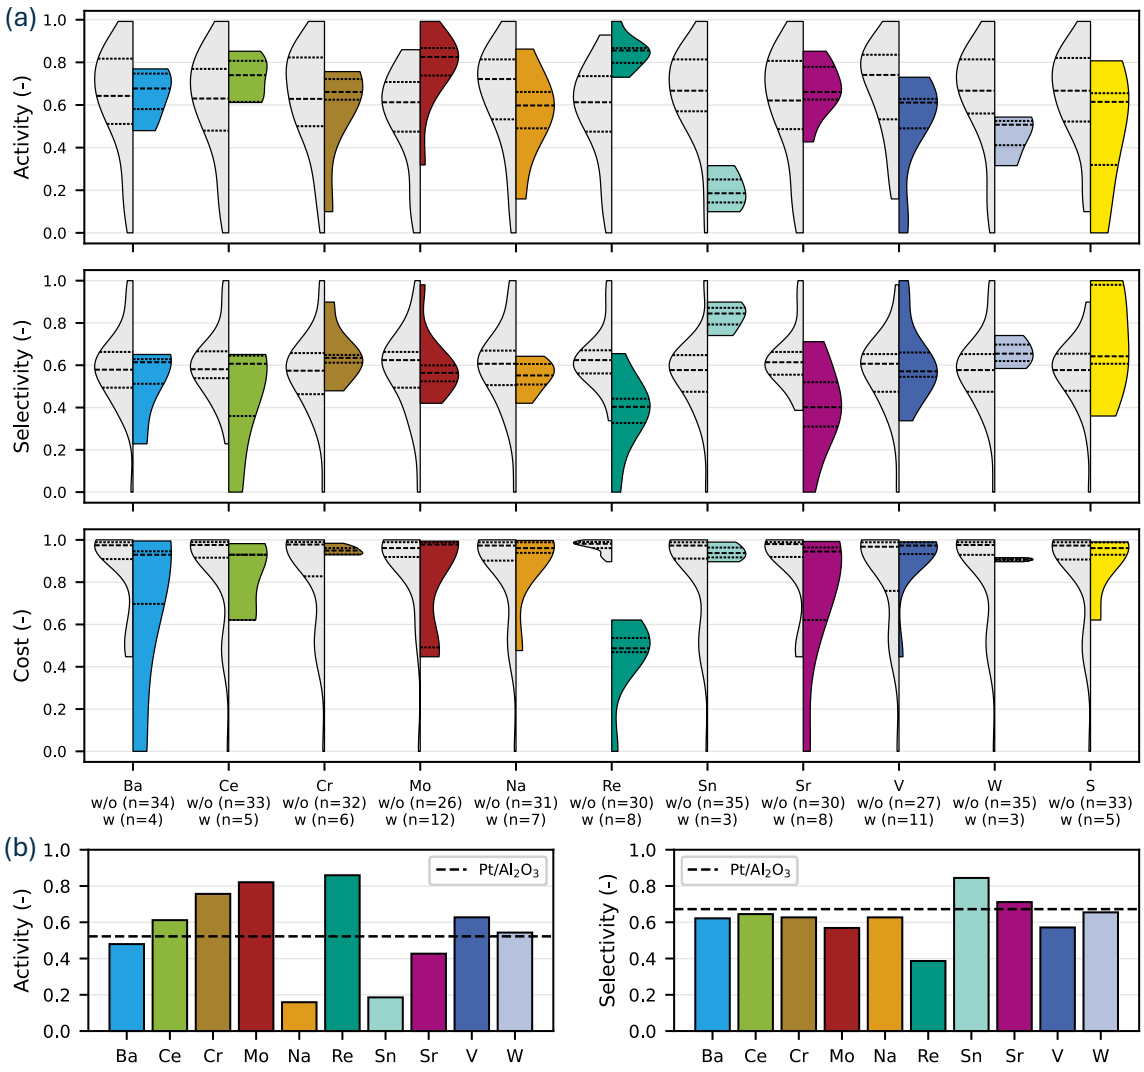


Figure 5: (a) Exploratory comparison of objectives for catalysts without (w/o) and with (w) individual promoter species within the design space. For each promoter, the left side of the violin represents catalysts not containing the promoter, while the right side represents catalysts containing the promoter. Horizontal lines denote the median and interquartile range. (b) Isolated effect of individual promoter species on catalyst activity and selectivity relative to the unpromoted Pt/Al_2O_3 reference catalyst (dashed line). All data points correspond to catalysts containing Pt and a single additional promoter (w_X = 0.30 wt.%, w_{Re} = 0.15 wt.%) under otherwise identical conditions.

When used as sole promoter, V and Cr show a small activity enhancing effect, in agreement with observations reported for the dehydrogenation of other hydrocarbons. A similar effect is observed for Ce promotion. In contrast, Na strongly suppresses catalytic activity, consistent with activity inhibiting effects reported in a different catalytic system. A less strong activity-reducing effect is observed for Sr. Ba and W promotion shows only marginal effects on both activity and selectivity in the isolated comparisons, indicating a largely neutral influence under the investigated conditions.

Although no isolated data point is available to directly quantify the effect of S promotion, both the exploratory analysis and literature data consistently indicate a positive effect on the selectivity. More generally, discrepancies between trends observed in the isolated and multivariate comparisons highlight the composition-dependent nature of promoter

effects, which arise from interactions between multiple promoter species and variations in promoter loadings within the broader catalyst design space.

3.2. Assessment of GP Kernel

To establish a suitable surrogate model for the subsequent catalyst optimisation, the predictive performance of different GP kernels (Matérn, RBF, and RQ) was systematically evaluated. Across all three objectives, the predictive performance of the tested GP kernels is mostly comparable (Figure 6a). For activity and selectivity, mean absolute errors (MAE) and associated standard deviations differ only moderately between kernels and validation strategies, indicating that predictive performance is primarily limited by data complexity rather than kernel choice. Regarding the different cross-validation strategies, all kernels exhibit the same consistent trend (LOOCV < random split 85% < random split 50%) in terms of prediction error and standard deviation. LOOCV leverages nearly the full dataset for training in each fold, resulting in lower prediction errors and reduced variance when compared with the smaller training splits, whereas smaller training sets increase both the error and its variance.

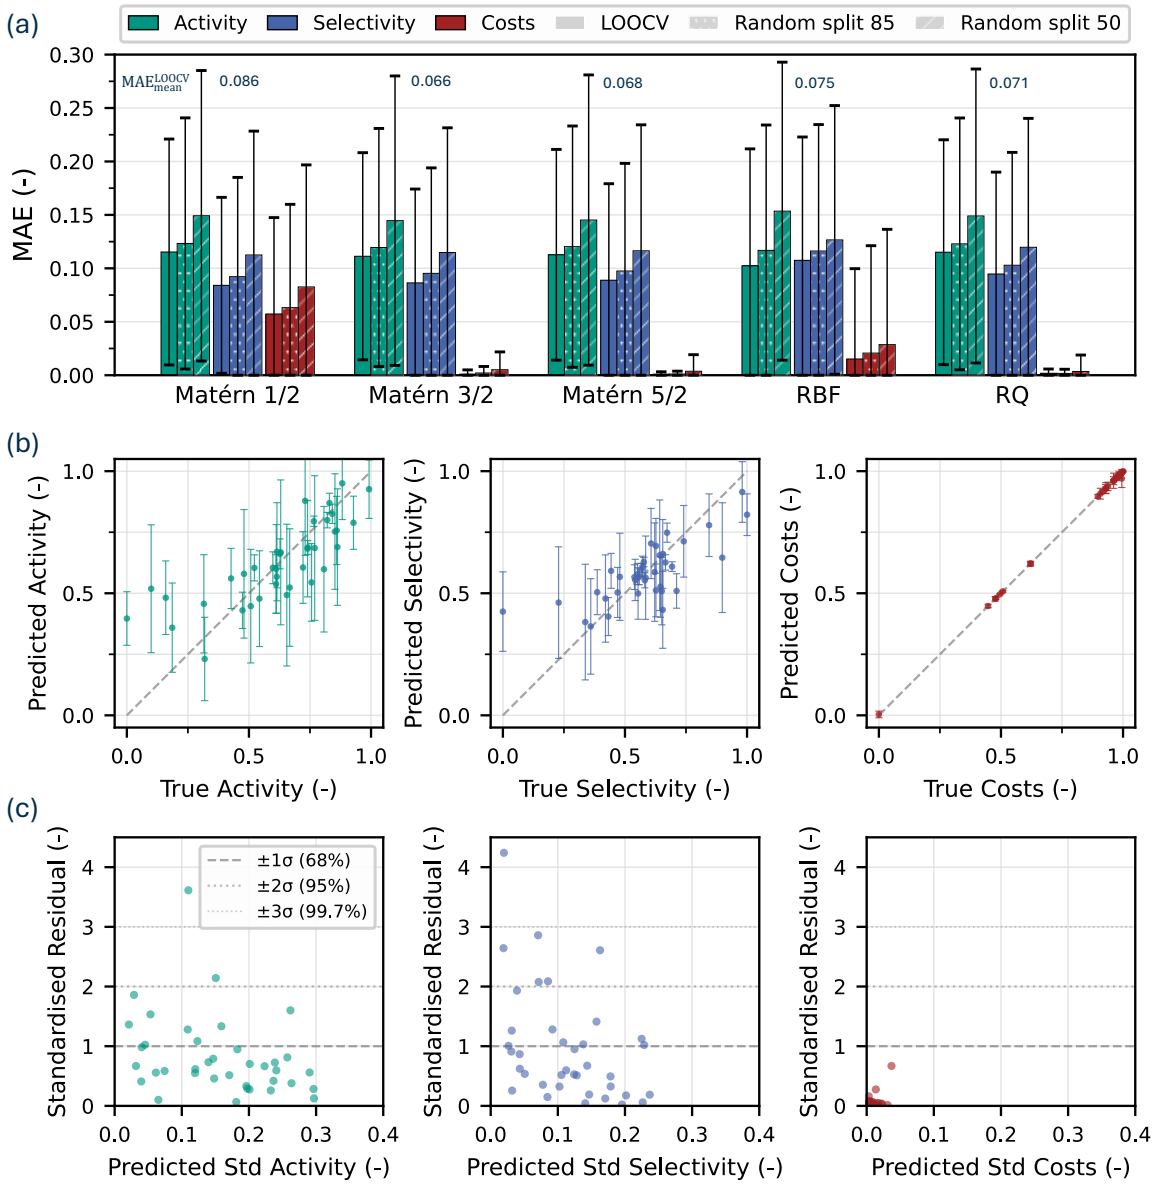


Figure 6: Comparison of GP kernels used as surrogate models. (a) Mean absolute error (MAE) with standard deviation obtained from leave-one-out cross-validation (LOOCV) and random train–test splits (85/15 and 50/50) for all three objectives. (b) Parity plots for LOOCV predictions with uncertainty across the three objectives with Matérn 3/2 kernel. (c) Calibration plots showing the standardised residuals (defined as the prediction error normalised by the predicted standard deviation) as a function of the predicted standard deviation for each objective with Matérn 3/2 kernel.

A clearer distinction emerges between objectives. Cost shows substantially lower prediction errors across all kernels, which is in line with a smooth and largely deterministic dependence on catalyst composition and precursor prices. In contrast, activity and selectivity are governed by complex physicochemical interactions, resulting in higher uncertainty. The use of independent surrogate models further isolates the well-behaved cost objective from potential negative transfer effects.

Catalyst-specific prediction errors further reveal substantial variation in predictive accuracy across the dataset (see section S1.4.). While many catalysts are predicted with relatively low absolute activity and selectivity errors <0.1, some candidates show

substantially larger errors, reaching approximately 0.40–0.45. These deviations may reflect combined effects of promoter identity and loading, number of promoters, as well as the limited number of experimental data points available to represent the 11-dimensional composition space.

Given the overall similar performance, the Matérn 3/2 kernel was selected for subsequent modelling as a balanced compromise between flexibility and regularity, providing consistently low error without imposing overly strong smoothness assumptions.

For activity and selectivity, predicted values in LOOCV closely follow the parity line, with no strong systematic bias across the target range (Figure 6b). Prediction uncertainty increases in regions with sparse data coverage, as reflected by larger deviations and error bars, which indicates that the surrogate appropriately reflects epistemic uncertainty.

Normalised prediction errors reveal well-calibrated uncertainty estimates for all three objectives, with prediction uncertainty spanning a meaningful range and most residuals remaining within the standard deviation (Figure 6c). For activity, approximately two thirds of the standardised residuals fall within one predicted standard deviation, with only few points exceeding the two-standard-deviation threshold, in line with expectations for a well-calibrated Gaussian predictive distribution. In contrast, uncertainty for costs predictions collapses across the design space, due to aforementioned deterministic behaviour. Selectivity shows a slightly heavier tail, including isolated deviations beyond two and three standard deviations, indicating localised overconfidence of the surrogate.

3.3. Influence of the Acquisition Function

Having identified a suitable surrogate model, the impact of different acquisition strategies on the optimisation performance is assessed. In the SO activity setting (Figure 7a), all evaluated acquisition functions (EI, NEI, and UCB) significantly outperform random search and converge rapidly toward the global optimum. EI, NEI and UCB exhibit similar convergence behaviour and reach the optimum within approximately 12–15 selection steps.

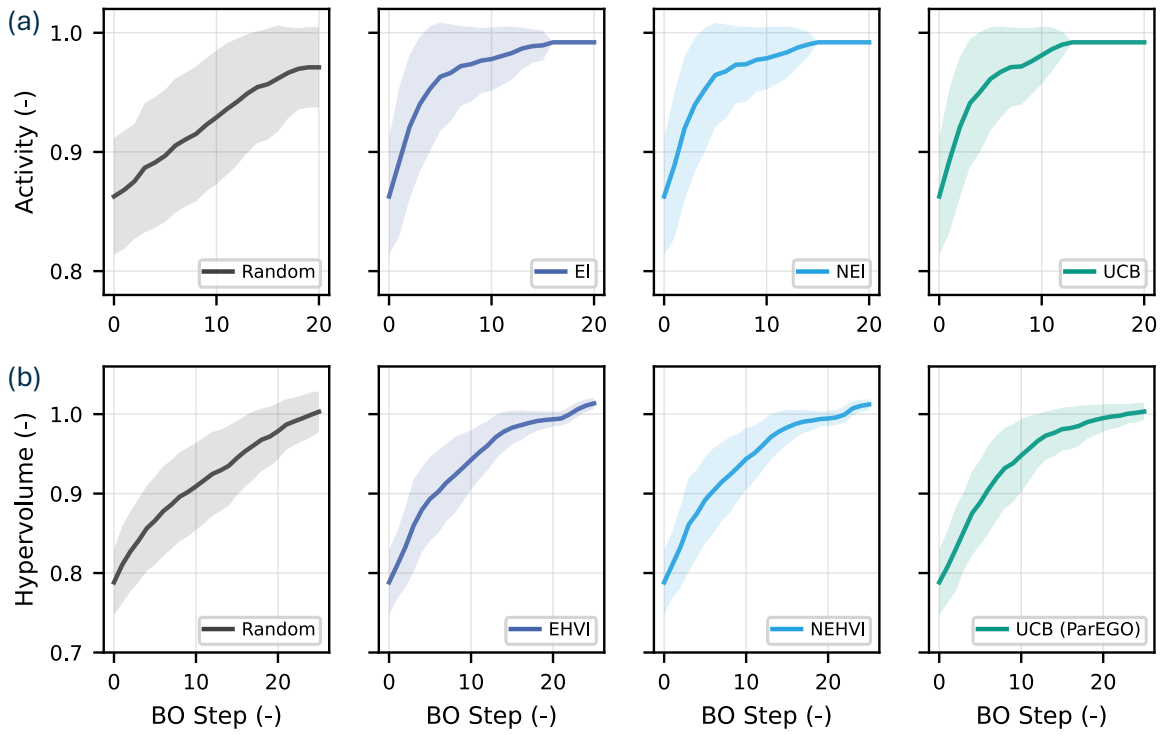


Figure 7: Retrospective analysis of acquisition function behaviour over successive BO selection steps. (a) Single-objective optimisation (activity) evaluated using the best observed score as performance metric. (b) Multi-objective optimisation evaluated using the hypervolume (HV) indicator. Shaded regions indicate the standard deviation across 100 runs.

In the MO setting (Figure 7b), optimisation performance is assessed *via* the hypervolume indicator. All acquisition functions achieve a monotonic increase in hypervolume and substantially outperform random sampling. The observed difference between EHVI, NEHVI, and UCB is minimal, which, once again, suggests that data complexity dominates performance under the given evaluation budget rather than acquisition choice. The monotonic increase in hypervolume reflects the progressive refinement of the Pareto front during optimisation. Overall, the results demonstrate that BO reliably accelerates improvement in both SO and MO settings, with particularly pronounced gains over random search. Due to comparable performances across acquisition functions, we employed NEI and NEHVI for subsequent catalyst optimisation due to their robustness to observation noise and their widespread use in recent BO literature.[34]

3.4. Data-driven Catalyst Optimisation

Both SO and MO optimisation scenarios were investigated to illustrate different use cases of BO for catalyst design. SO optimisation of catalyst costs is uninformative, as the unpromoted 0.30Pt catalyst already represents the cost-minimum solution within the defined design space. Likewise, optimisation targeting selectivity exclusively is of limited practical relevance as high S loadings effectively suppress MeFl formation by deactivating the catalyst.

In contrast, activity-driven SO is a useful optimisation, as catalytic activity cannot simply be increased by increasing the loading of promoters, such as Re.[13] The catalysts evaluated in the three SO optimisation iterations performed in this study are consistently ranked among the most active candidates (Figure 8a), which points towards effective prioritisation of promising candidates in the optimisation framework. Moreover, the experimentally measured activities closely follow the predicted values within the associated uncertainty bounds, i.e. the surrogate model provides a meaningful ranking of promising candidates.

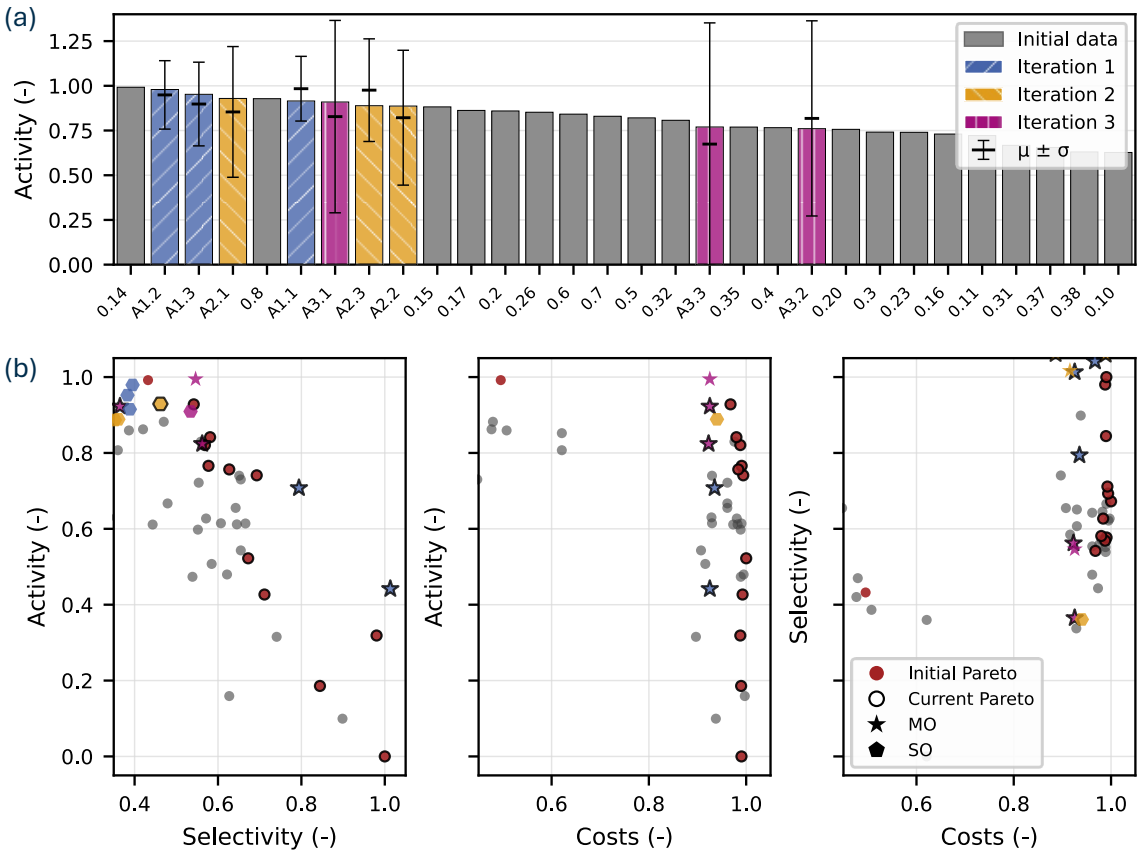


Figure 8: (a) Ranking of the top 30 catalysts based on measured activity. Catalysts evaluated in the three optimisation iterations of the single-objective (SO) activity-driven optimisation are highlighted, together with the corresponding predicted mean and uncertainty from the GP posterior. (b) Two-dimensional projections of the three-dimensional objective space with the empirical Pareto front of the initial dataset indicated. Catalysts evaluated during the three optimisation iterations in both the SO and multi-objective (MO) settings are overlaid.

While the SO case allows for a straightforward assessment of optimisation progress based on ranking, the MO case requires a different perspective. Here, progress is assessed based on the relative position of newly evaluated catalysts with respect to the empirical Pareto front defined by $\mathcal{D}_{\text{start}}$. The catalysts proposed in optimisation iterations 1-3 populate regions of the objective space that are sparsely covered by the initial dataset and lie close to the previously observed Pareto front or even manage to expand the Pareto front (Figure 8b). This behaviour indicates that BO effectively explores the Pareto front by identifying diverse trade-off solutions between activity, selectivity, and costs. Consequently, some proposed candidates exhibit comparatively low activity while

offering favourable performance with respect to the remaining objectives. Some of the proposed MO candidates lie on the Pareto front despite exhibiting a strongly reduced activity <0. While some of these trade-off solutions may not represent practically desirable catalyst compositions, they naturally arise in unconstrained multi-objective optimisation. In future studies, additional constraints, such as a minimum activity threshold, could be incorporated to restrict the optimisation to application-relevant regions of the design space.

The evolution of catalyst compositions during the optimisation iterations reveals both consistent trends and mode-specific differences between the SO and MO settings. In both optimisation modes, Mo is proposed as promoter for most catalysts. This highlights the great potential as promoter within the explored compositional space (Figure 9).

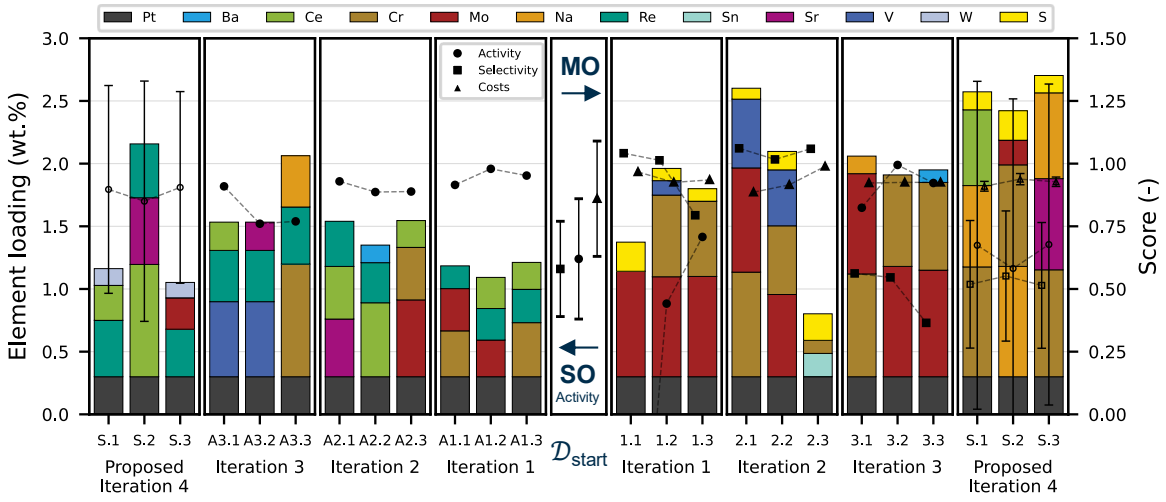


Figure 9: Centre: training data with mean activity, selectivity and cost and each standard deviation. To the left and right are the single objective (SO) and multi objective (MO) optimisation iterations, respectively, showing the composition and measured scores of experimentally evaluated catalysts. To the very left and right suggestions for the next SO and MO iteration, respectively, with predicted Scores and uncertainties from the GP posterior. Dashed lines for better visualisation purpose only.

In the SO activity-driven optimisation, the first three iterations already yield catalysts exhibiting activity values exceeding the mean activity of the initial dataset by more than 66%. These catalysts are characterised by moderate total promoter loadings in the range of approximately 1…1.75 wt.%. Frequently observed promoters include Re, Cr and Ce. Although certain promoter combinations occur repeatedly, both evaluated catalysts and newly proposed candidates frequently include additional promoter species, reflecting the current balance between exploration of sparsely sampled regions of the design space and exploitation of promising activity trends.

In contrast, the MO optimisation exhibits a distinct compositional focus with higher overall promotor loading. Many proposed candidates contain S, which is consistent with its well-known role in suppressing undesired side reactions. To compensate for the associated loss in activity, S is combined with activity-enhancing promoters such as Mo and Cr. This compositional balancing illustrates how the MO framework explicitly

navigates trade-offs between activity, selectivity, and cost rather than optimising a single performance metric.

The comparatively high predictive uncertainty of the proposed candidates for Iteration 4 further indicates that the optimisation is still in an exploratory phase, which is expected given the large combinatorial design space and the limited amount of experimental data currently available. As additional data are acquired, the surrogate model is expected to become more confident, enabling increasingly exploitative candidate selection and further refinement of promising catalyst compositions.

4. Conclusion and outlook

Overall, this work demonstrates how BO can be practically employed as a data-efficient decision-support tool for catalyst design in a case study for LOHC dehydrogenation. Rather than aiming at algorithmic benchmarking or identifying a globally optimal catalyst composition, the focus lies on demonstrating guidance of BO through experimental exploration and prioritisation of promising catalyst candidates under realistic laboratory constraints. In addition, the systematic analysis of the available experimental data provides insights into the influence of individual promoters and compositional trends, complementing the optimisation framework with a better understanding of catalyst behaviour. Despite the limited number of optimisation iterations, the presented framework illustrates the potential of BO to efficiently navigate large catalyst design spaces and manages to expand the Pareto front in the MO setting.

Future work will extend both the optimisation campaigns and the BO framework itself. Larger experimental datasets and a substantially increased number of iterations are expected to improve surrogate model accuracy and enable a gradual shift from exploration towards more exploitative candidate selection. In parallel, the catalyst design space will be expanded to include additional promoter elements, alternative support materials, and varying Pt loadings, allowing a more comprehensive exploration of complex catalyst systems. Beyond expanding the catalyst composition, future studies may also investigate the simultaneous optimisation of synthesis parameters, such as calcination and reduction conditions. This would allow for the optimisation framework to account for promoter-specific preparation requirements and may further improve catalyst performance.

Beyond the present demonstration study, the optimisation framework itself can be further refined by investigating alternative surrogate models, acquisition strategies, and optimisation settings tailored to catalyst development. In particular, compositional constraints such as the limit on the number of promoters (Section 2.3) could not be enforced directly during acquisition and instead relied on pre-generating a discrete candidate pool via Sobol sampling within the constrained design space. As the design space grows, this approach becomes increasingly limiting, since only a finite, precomputed set of compositions is ever considered. A promising direction to address

this is coupling GP surrogate modelling with a synthesis-aware genetic algorithm, which would allow such constraints to be enforced directly during candidate generation rather than through discrete pre-filtering.

Acknowledgements

Financial support by the Helmholtz Research Program “Materials and Technologies for the Energy Transition (MTET), Topic 3: Chemical Energy Carriers” and “Innovationspool – Auf dem Weg zur autonomen Materialentwicklung” is highly acknowledged. Further, financial support by the Erasmus+ programme of the European Union for Torben Tiezema is acknowledged. Laura Marie Helleckes is supported by an Eric and Wendy Schmidt AI in Science Postdoctoral Fellowship, a Schmidt Sciences program. The authors gratefully acknowledge Friedolin Hamperl for contributing to catalyst synthesis and conducting catalytic experiments and Armin Lautenbach for ICP-AES measurements.

References

1 T. Rüde, S. Dürr, P. Preuster, M. Wolf and P. Wasserscheid, Benzyltoluene/perhydro benzyltoluene – pushing the performance limits of pure hydrocarbon liquid organic hydrogen carrier (LOHC) systems, *Sustain. Energy Fuels*, 2022, **6**, 1541–1553, DOI: 10.1039/D1SE01767E.

2 M. R. Usman, D. L. Cresswell and A. A. Garforth, By-Products Formation in the Dehydrogenation of Methylcyclohexane, *Pet. Sci. Technol.*, 2011, **29**, 2247–2257, DOI: 10.1080/10916466.2011.584103.

3 P. Preuster, A. Alekseev and P. Wasserscheid, Hydrogen Storage Technologies for Future Energy Systems, *Annu. Rev. Chem. Biomol. Eng.*, 2017, **8**, 445–471, DOI: 10.1146/annurev-chembioeng-060816-101334.

4 J. Humphreys, R. Lan and S. Tao, Development and Recent Progress on Ammonia Synthesis Catalysts for Haber–Bosch Process, *Adv. Energy Sustain. Res.*, 2021, **2**, 2000043, DOI: 10.1002/aesr.202000043.

5 Q. Zhang, W. Deng and Y. Wang, Recent advances in understanding the key catalyst factors for Fischer-Tropsch synthesis, *J. Energy Chem.*, 2013, **22**, 27–38, DOI: 10.1016/S2095-4956(13)60003-0.

6 U. Kurumbail, H. R. Darji, M. Alvear, S. Chen and I. Hermans, A case study in the development of improved promoted Pt catalysts for propane dehydrogenation through Bayesian optimization with uncertainty quantification, *Chem. Eng. J.*, 2025, **505**, 158927, DOI: 10.1016/j.cej.2024.158927.

7 R. A. Van Santen, Chemical basis of metal catalyst promotion, *Surf. Sci.*, 1991, **251–252**, 6–11, DOI: 10.1016/0039-6028(91)90944-N.

8 Wasserscheid Peter, Boesmann Andreas, Seidel Alexander, Frankl Franziska, Sturm Stefanie, and Szesni Normen, WO-2019201706-A1 -Platinum-sulfur-based Shell Catalyst, Production and Use Thereof in the Dehydrogenation of Hydrocarbons.

9 K. Fukudome, K. Imagawa, H. Kobayashi, Y. Okada and M. Saito, Egg shell-type platinum-loaded alumina catalyst, method of producing same, and method of using same, AU2020443246A1, 2023.

10 A. D. Ballarini, P. Zgolicz, I. M. J. Vilella, S. R. De Miguel, A. A. Castro and O. A. Scelza, n-Butane dehydrogenation on Pt, PtSn and PtGe supported on γ-Al_2O_3 deposited on

spheres of α-Al_2O_3 by washcoating, *Appl. Catal. Gen.*, 2010, **381**, 83–91, DOI: 10.1016/j.apcata.2010.03.053.
11 K. Alconada and V. L. Barrio, Enhancing perhydrobenzyltoluene dehydrogenation performance with Co, Mo and Mn metal oxides: A comparative study with Pt/Al_2O_3 catalyst, *Appl. Catal. B Environ. Energy*, 2024, **357**, 124349, DOI: 10.1016/j.apcatb.2024.124349.
12 K. Alconada and V. L. Barrio, Evaluation of bimetallic Pt–Co and Pt–Ni catalysts in LOHC dehydrogenation, *Int. J. Hydrog. Energy*, 2024, **51**, 243–255, DOI: 10.1016/j.ijhydene.2023.04.120.
13 D. Strauch, P. Weiner, B. B. Sarma, A. Körner, E. Herzinger, P. Wolf, A. Zimina, A. Hutzler, D. E. Doronkin, J.-D. Grunwaldt, P. Wasserscheid and M. Wolf, Bimetallic platinum rhenium catalyst for efficient low temperature dehydrogenation of perhydro benzyltoluene, *Catal. Sci. Technol.*, 2024, **14**, 1775–1790, DOI: 10.1039/D3CY01336G.
14 B. Bong, C. Mebrahtu, D. Jurado, A. Bösmann, P. Wasserscheid and R. Palkovits, Hydrogen Loading and Release Potential of the LOHC System Benzyltoluene/Perhydro Benzyltoluene over S–Pt/TiO_2 Catalyst, *ACS Eng. Au*, 2024, **4**, 359–367, DOI: 10.1021/acsengineeringau.4c00003.
15 F. Auer, D. Blaumeiser, T. Bauer, A. Bösmann, N. Szesni, J. Libuda and P. Wasserscheid, Boosting the activity of hydrogen release from liquid organic hydrogen carrier systems by sulfur-additives to Pt on alumina catalysts, *Catal. Sci. Technol.*, 2019, **9**, 3537–3547, DOI: 10.1039/C9CY00817A.
16 A. D. Ballarini, S. R. De Miguel, A. A. Castro and O. A. Scelza, n-Decane dehydrogenation on Pt, PtSn and PtGe supported on spinels prepared by different methods of synthesis, *Appl. Catal. Gen.*, 2013, **467**, 235–245, DOI: 10.1016/j.apcata.2013.07.030.
17 M.-H. Lee, B. M. Nagaraja, K. Y. Lee and K.-D. Jung, Dehydrogenation of alkane to light olefin over PtSn/θ-Al_2O_3 catalyst: Effects of Sn loading, *Catal. Today*, 2014, **232**, 53–62, DOI: 10.1016/j.cattod.2013.10.011.
18 Y. Zhang, Y. Zhou, A. Qiu, Y. Wang, Y. Xu and P. Wu, Propane dehydrogenation on PtSn/ZSM-5 catalyst: Effect of tin as a promoter, *Catal. Commun.*, 2006, **7**, 860–866, DOI: 10.1016/j.catcom.2006.03.016.
19 S. C. Purdy, P. Ghanekar, G. Mitchell, A. J. Kropf, D. Y. Zemlyanov, Y. Ren, F. Ribeiro, W. N. Delgass, J. Greeley and J. T. Miller, Origin of Electronic Modification of Platinum in a Pt_3 V Alloy and Its Consequences for Propane Dehydrogenation Catalysis, *ACS Appl. Energy Mater.*, 2020, **3**, 1410–1422, DOI: 10.1021/acsaem.9b01373.
20 Y. Zhou, S. Qi, B. Yang and C. Yi, WPt/Al_2O_3 Catalyst for 18H-DBT Long-Period Dehydrogenation in a Fixed Bed Reactor at a Low Temperature, *Ind. Eng. Chem. Res.*, 2025, **64**, 20961–20970, DOI: 10.1021/acs.iecr.5c03041.
21 P. Courtine and E. Bordes, *Synergistic effects in multicomponent catalysts for selective oxidation*, Elsevier, 1997, vol. 110.
22 M. Caldas Ramos, S. S. Michtavy, A. D. White and M. D. Porosoff, Bayesian Optimization of Catalysis with In-Context Learning, *ACS Cent. Sci.*, 2026, **12**, 599–615, DOI: 10.1021/acscentsci.5c02418.
23 A. Ramirez, E. Lam, D. P. Gutierrez, Y. Hou, H. Tribukait, L. M. Roch, C. Copéret and P. Laveille, Accelerated exploration of heterogeneous CO_2 hydrogenation catalysts by Bayesian-optimized high-throughput and automated experimentation, *Chem Catal.*, 2024, **4**, 100888, DOI: 10.1016/j.checat.2023.100888.

24J. L. Santos, D. K. Alsaadi, A. Alahmadi, D. Khizbullin, D. Pugh, A. Genovese, M. H. Alabsi, K. A. Al-Majnouni, M. Khawaji, J. M. R. Gallo and J. Gascon, Bayesian optimization guides the discovery of high-performance catalysts for direct CO_2-to-jet-fuel conversion, *Chem Catal.*, 2026, **6**, 101683, DOI: 10.1016/j.checat.2026.101683.
25J. K. Pedersen, C. M. Clausen, O. A. Krysiak, B. Xiao, T. A. A. Batchelor, T. Löffler, V. A. Mints, L. Banko, M. Arenz, A. Savan, W. Schuhmann, A. Ludwig and J. Rossmeisl, Bayesian Optimization of High-Entropy Alloy Compositions for Electrocatalytic Oxygen Reduction, *arXiv*, 2021, preprint, DOI: 10.48550/ARXIV.2106.08212.
26G. O. Kayode, A. F. Hill and M. M. Montemore, Bayesian optimization of single-atom alloys and other bimetallics: efficient screening for alkane transformations, CO_2 reduction, and hydrogen evolution, *J. Mater. Chem. A*, 2023, **11**, 19128–19137, DOI: 10.1039/D3TA02830E.
27F. Thelen, R. Zehl, R. Zerdoumi, J. L. Bürgel, L. Banko, W. Schuhmann and A. Ludwig, Accelerating Combinatorial Electrocatalyst Discovery with Bayesian Optimization: A Case Study in the Quaternary System Ni-Pd-Pt-Ru for the Oxygen Evolution Reaction, *arXiv*, 2025, preprint, DOI: 10.48550/ARXIV.2502.16241.
28C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*, MIT Press, 2006.
29D. R. Jones, M. Schonlau and W. J. Welch, Efficient Global Optimization of Expensive Black-Box Functions, *J. Glob. Optim.*, 1998, DOI: 10.1023/A:1008306431147.
30C. Robert and G. Casella, *Monte Carlo Statistical Methods*, Springer, 2004.
31I. M. Sobol, On the Distribution of Points in a Cube, *USSR Comput. Math. Math. Phys.*, 1967.
32N. Srinivas, A. Krause, S. M. Kakade and M. Seeger, Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design, *ICML*, 2010, DOI: https://doi.org/10.1109/TIT.2011.2182033.
33R. Garnett, *Bayesian Optimization*, Cambridge University Press, Cambridge, 2023.
34M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson and E. Bakshy, BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization, *Adv. Neural Inf. Process. Syst. 33*, 2020, DOI: https://doi.org/10.48550/arXiv.1910.06403.
35A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai and S. Chintala, PyTorch: An Imperative Style, High-Performance Deep Learning Library, *arXiv*, 2019, preprint, DOI: 10.48550/ARXIV.1912.01703.
36J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger and A. G. Wilson, GPyTorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration, *arXiv*, 2018, preprint, DOI: 10.48550/ARXIV.1809.11165.
37B. Letham, B. Karrer, G. Ottoni and E. Bakshy, Constrained Bayesian Optimization with Noisy Experiments, *NeurIPS*, 2019, DOI: https://doi.org/10.48550/arXiv.1706.07094.
38M. Emmerich, K. C. Giannakoglou and B. Naujols, Single- and Multiobjective Optimization Assisted by Gaussian Random Field Metamodels, *IEEE Trans. Evol. Comput.*, 2006, DOI: 10.1109/TEVC.2005.859463.
39S. Daulton, M. Balandat and E. Bakshy, Differentiable Expected Hypervolume Improvement for Parallel Multi-Objective Bayesian Optimization, *NeurIPS*, 2021, DOI: https://doi.org/10.48550/arXiv.2006.05078.

40J. Knowles, ParEGO: A Hybrid Algorithm with On-line Landscape Approximation for Expensive Multiobjective Optimization Problems, *Evol. Comput.*, 2006, DOI: 10.1109/TEVC.2005.851274.